Predicting Central Topics in a Blog Corpus from a Networks Perspective

Srayan Datta
University of Michigan
Ann Arbor, Michigan 48105
srayand@umich.edu

arXiv:1405.2386v1 [cs.IR] 10 May 2014

Abstract

In today's content-centric Internet, blogs are becoming increasingly popular and important from a data analysis perspective. According to Wikipedia, there were over 156 million public blogs on the Internet as of February 2011. Blogs are a reflection of our contemporary society. The contents of different blog posts are important from social, psychological, economical and political perspectives. Discovery of important topics in the blogosphere is an area which still needs much exploring. We try to come up with a procedure using probabilistic topic modeling and network centrality measures which identifies the central topics in a blog corpus.

1 Introduction

This paper presents an algorithm to identify and rank the most important topics given a set of blog entries. The topics are identified using Latent Dirichlet Allocation (LDA) (Blei et al., 2003) – a probabilistic topic model, and subsequently labeled to indicate their identity. Topic networks are created from the topics extracted using two different metrics. Graph centrality measures are run on these networks to find out two different rankings. A topic navigator is built to display the most relevant documents given a topic, where the topics are ranked by their importance given a similarity metric.

2 Design

This section is divided into several parts. The first sub-section describes the dataset. The next one discusses data preprocessing and topic extraction, which is followed by the subsections topic network creation, centrality calculation, manual labeling and topic navigator description.

2.1 Dataset description

Due to crawling being extremely time-consuming, a suitable publicly available blog corpus called the Blog Authorship Corpus (Schler et al., 2006), is used for this project. This corpus contains around 681,000 posts from blogger.com. All of these blogs contain authorship information (to filter out some of the junk posts) and there are at least 200 words in each blog. The blog corpus contains 19,320 XML files where each file stores all the blog posts of a specific author.

2.2 Data preprocessing and topic extraction

The topic extraction is performed using LDA topic model, a very popular choice among different probabilistic topic models. The topic extraction is performed using a popular topic modeling toolkit called Mallet (McCallum, 2002). As Mallet does not take XML files as input, each XML file is converted to a text file (.txt format) by stripping the XML metadata. Presence of stop words significantly deteriorates topic model performance, so stop word removal is done using Mallet.

Small text volume in a document affects LDA performance. Stop words represent a significant volume of any text. After stop word removal, the number of words in a blog entry decreases even more. All the blog entries by a single author are stored into a single file to alleviate this problem.

Finding optimal number of topics given a corpus is a difficult problem. For the purpose of this project this number is assumed to be 100. Mallet outputs the most important 19 words for each topic and the probability distribution of the 100 topics over each file.

2.3 Manual labeling

As Mallet does not output the names of the topics, the top 19 words for each topic emitted by Mallet are used to manually label the topics. This is done so that each topic has an identity and is not just a

bag of words. Assigning a name from just a list of words is a challenging task. In this case, this is even more challenging, as in the blog corpus, the language is very informal and there are languages other than English. There are cases where some proper nouns (names of persons) are returned as the most important words. These names do not give any idea about the underlying topic. Mallet is seen to output some mixed topics also. There are two topics for which no suitable nomenclature is found. These two topics are denoted by 'Mixed'.

2.4 Topic network creation

In this section, the methods of building topic networks from Mallet are discussed. A topic network is a weighted graph where each node is specified by a topic and an weighted edge between two nodes denotes the degree of similarity between the topics of those two nodes. Note that an edge weight is always positive.

Building a topic network can be divided into two parts; identifying features for each topic and deciding which similarity measure works best. For the first part, we have two different feature vectors for each topic. Mallet outputs the probability of a word given a topic for each topic and each word (henceforth referred to as *topic-word features*); and probability of a topic given a document for each document and each topic (henceforth referred to as *topic-document features*). The topic-word feature becomes intractable due to a very high number of unique words in the dictionary. This number is very high as all derivatives of a single word are considered unique and there are words from languages other than English in the blog corpus. So, the topic-document features are used instead. Each topic t_i is represented by a feature vector $< p_{ij} >$ of length 19,320 (number of documents in the blog corpus), where p_{ij} is the probability that the j^{th} document d_j belongs to the i^{th} topic t_i. These feature vectors are then used to calculate similarity between two topics.

Due to a lack of previous work of this nature, it is difficult to tell which similarity measure between the feature vectors will work the best in this case. So, instead of choosing a single measure, two very popular similarity measures were used to create two topic networks. The first one is cosine similarity (Wikipedia, 2013a), which is a widely used metric in information retrieval and natural language processing, where it is used to measure similarity between documents, words, sentences, and so on. The other metric used is Pearson correlation coefficient (Wikipedia, 2013b) which is a measure of linear dependence between two random variables in statistics.

Cosine similarity always lies between zero and one. So, when using this measure to build a topic network, there is always an edge with positive weight between any two nodes. So, the graph generated is a complete graph. On the other hand, correlation coefficient can take positive, negative and zero values. The negative and zero weight edges are pruned due to the fact that there is no well-known centrality measure which works well on a weighted graph with negative edge weights. So, the graph generated using correlation coefficient as a similarity measure is not a complete graph.

2.5 Centrality calculation

There are quite a few ways of assigning importance to nodes in a weighted graph using only graph structure. In this case we use the most famous one, Pagerank (Brin and Page, 1998), which is shown to have good performance over a range of different networks including Internet link graph, citation graphs, social networks and word networks (Mihalcea and Tarau, 2004). Pagerank is run on both topic networks and the topics are ranked according to their Pagerank values.

Other popular centrality measures like degree and betweenness does not perform well because one of the topic networks is a complete graph.

2.6 Topic navigator description

A topic navigator augments the work presented above. The navigator provides a simple graphical user interface which allows a user to select a similarity metric from a drop-down list. Another drop-down menu showing the ranked list of topics according to the selected similarity measure is shown. This list is scrollable and selecting a topic shows the top 100 blogs related to the topic, ranked from most relevant to least relevant. Selecting a blog allows scrolling; double-clicking opens the document in the default text editor. Reselecting a metric and a topic dynamically changes the results shown.

Figure 1 shows a screen-shot of the topic navigator.

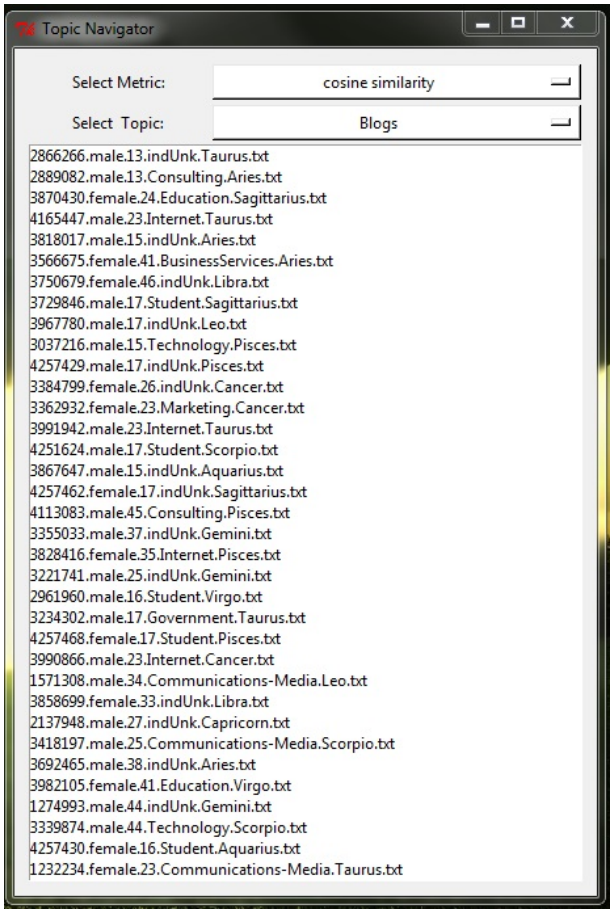

Figure 1: The Topic Navigator

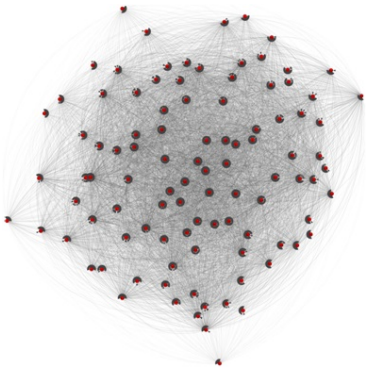

Figure 2: The Topic Network using cosine similarity

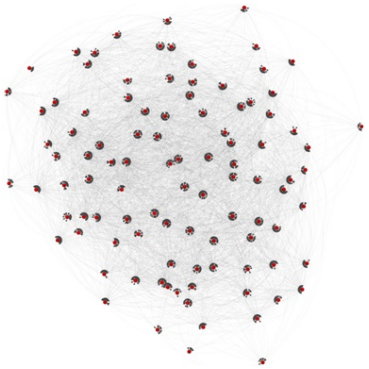

Figure 3: The Topic Network using correlation coefficient

3 Implementation and tools used

The Topic Modelling is done using Mallet (McCallum, 2002), as mentioned before. The rest of the implementation is done using python. Igraph library (Csardi and Nepusz, 2006) for python is used for pagerank implementation. The graph visualizations are generated using igraph and Gephi (Bastian et al., 2009).

4 Results

This section presents the results obtained from the experiments described above. Figure 2 presents a visualization of the topic network created using cosine similarity and figure 3 presents visualization of the topic network using correlation coefficient. These graphs give a broad idea about the nature of connection between the topics and the difference between the two topic networks.

We present top 10 topics determined by each similarity metric in a tabular format. The top 19 words for each topic are also presented. Table 1 shows the top topics for cosine similarity measure and Table 2 likewise shows the same for correlation coefficient. Note that, some topics occur in both tables.

We present visualizations of the ranked topics in the form of 'topic cloud's . A topic cloud is a visual representation of topics in a corpus where the font size and color of the topic are determined by its importance. The larger and darker the topic, the more its importance. Note that unlike tag clouds (generally used for free-from text visualization), topic clouds maintain their underlying network structure. So, if two topics are nearby, they are more closely related to each other than two topics that lie at the two extreme ends of the picture. Figure 4 shows the topic cloud derived from cosine similarity topic network. Likewise, figure 5 represents the topic cloud using correlation coefficient.

From the cosine similarity table, it is interesting to observe that the topics like 'Idle chitchat' and 'Talking about people' are ranked very high. Moreover, 'Talking about people' is ranked high in correlation coefficient table. This, and the use of informal language, shows that a high number

Rank	Topic name	Top words
1	Idle chitchat	ve ll don didn pretty time thing back things good lot haven doesn wasn bit stuff guess isn won
2	Spent time	time fact year years point day long days made bit apparently half completely spent entire thing end case simply
3	Talking about people	don people didn doesn thing things isn make good bad thought wasn feel hate won wouldn person talking talk
4	Good times	good time great cool pretty fun day today stuff hope guys thing guess lot nice back work things people
5	Swear words normal	yeah hell damn ll crap thing hey gonna blah ass stuff kinda back god guy stupid man gotta head
6	Work life	people time things make work good person feel life find don lot point part thing place change interesting mind
7	Time description	today work day night time home morning back tomorrow good tonight week sleep bed days yesterday hours house weekend
8	Fashion trends	hair make big day wear wearing favorite shoes white hot black shirt put pants clothes red things buy kind
9	Life and death	people man world evil death kill men dead die make great life human hell thing years war fight live
10	How to write	ve words don word sense kind doesn real read high write sort writing isn makes line form piece sound

Table 1: Top topics determined from cosine similarity topic network

Rank	Topic name	Top words
1	SMS lingo 1	lol cuz gonna ya yea kinda wanna ppl omg today tho hey lil soo dunno stuff fun guys guess
2	German words	tyke ich und tina die der ist das phil cuz mit nicht purdy dr mich es evan ein von
3	SMS lingo 2	da ur dat sum wat coz jus im bout wit ppl goin day hav skool wen lol rite nite
4	Life and death	people man world evil death kill men dead die make great life human hell thing years war fight live
5	Sexual preference	women sex gay men woman marriage sexual man male relationship married female love life lesbian straight make partner support
6	Friends	alex jeremy mark hannah blog luke cindy day level fun time people lesson sammy school ted house cool love
7	Education	school class students college law student classes year semester summer university campus teaching paper study professor high graduate teacher
8	Today	alot day today ya thing bored back till house good time school bye kinda friends post life start friend
9	Good times	good time great cool pretty fun day today stuff hope guys thing guess lot nice back work things people
10	Talking about people	don people didn doesn thing things isn make good bad thought wasn feel hate won wouldn person talking talk

Table 2: Top topics determined from correlation coefficient topic network

of blogs are used very casually in a conversational style. The top three topics in the correlation coefficient table clearly show that the top-ranked topics are influenced by the use of informal language and presence of foreign languages. Some potential improvements over these issues are discussed in the discussion and future works section.

This study is the first of its kind. So, there is no

Figure 4: The Topic Cloud using cosine similarity

Figure 5: The Topic Cloud using correlation coefficient

parallel set of results on topic ranking to compare with.

5 Discussion and Future Work

Working with a blog corpus necessitates working with informal language, lots of abbreviations and wrong spellings. This makes topic extraction results look very different than the results reported in the literature (e.g., Table 2 of (Wang et al., 2007)). Further, the blog corpus used in this project contains languages other than English. This makes those other languages (e.g., German) themselves emerge as 'topics', rather than the actual topics that they encode. (A blog written in a different language outputs the language itself as its main topic irrespective of the content.) This happens because LDA assumes a topic is simply a probability distribution over words. Creating a special stop word list for the blog corpus and restricting the language to English may alleviate these problems.

Manual labeling from single words is difficult. Use of n-gram topic model (Wallach, 2005; Wang et al., 2007) can help with this issue, as it outputs phrases along with words, thereby potentially being more effective from a 'topic labeling' viewpoint.

6 Conclusion

This study is the first of its kind. Topic modelling outputs the topics of a given corpus, the methodology described here provides a way to rank them. This paper describes how to extract important topics from blogs, but the algorithm provided is not limited to a blog corpus, it is domain and language independent. This model can be used to extract important topics from any monolingual corpus.

Acknowledgements

I would like to thank Professor Sugih Jamin for the topic navigator idea, and Mr. Shibamouli Lahiri for discussions about the project idea and help with manual labeling of topics.

References

[Bastian et al.2009] Mathieu Bastian, Sebastien Heymann, and Mathieu Jacomy. 2009. Gephi: An Open Source Software for Exploring and Manipulating Networks.

[Blei et al.2003] David M. Blei, Andrew Y. Ng, and Michael I. Jordan. 2003. Latent Dirichlet Allocation. *Journal of Machine Learning Research*, 3:993–1022.

[Brin and Page1998] Sergey Brin and Lawrence Page. 1998. The Anatomy of a Large-scale Hypertextual Web Search Engine. *Comput. Netw. ISDN Syst.*, 30(1-7):107–117, April.

[Csardi and Nepusz2006] Gabor Csardi and Tamas Nepusz. 2006. The igraph software package for complex network research. *InterJournal*, Complex Systems:1695.

[McCallum2002] Andrew Kachites McCallum. 2002. MALLET: A Machine Learning for Language Toolkit. http://www.cs.umass.edu/ mccallum/mallet.

[Mihalcea and Tarau2004] Rada Mihalcea and Paul Tarau. 2004. TextRank: Bringing Order into Texts. In Dekang Lin and Dekai Wu, editors, *Proceedings of EMNLP 2004*, pages 404–411, Barcelona, Spain, July. Association for Computational Linguistics.

[Schler et al.2006] Jonathan Schler, Moshe Koppel, Shlomo Argamon, and James W. Pennebaker. 2006. Effects of Age and Gender on Blogging. In *AAAI Spring Symposium: Computational Approaches to Analyzing Weblogs*, pages 199–205.

[Wallach2005] Hanna M. Wallach. 2005. Topic modeling: beyond bag-of-words. In *NIPS 2005 Workshop on Bayesian Methods for Natural Language Processing*.

[Wang et al.2007] Xuerui Wang, Andrew McCallum, and Xing Wei. 2007. Topical n-grams: Phrase and topic discovery, with an application to information retrieval. In *Proceedings of the 2007 Seventh IEEE International Conference on Data Mining*, ICDM '07, pages 697–702, Washington, DC, USA. IEEE Computer Society.

[Wikipedia2013a] Wikipedia. 2013a. Cosine similarity — Wikipedia, the free encyclopedia. [Online; accessed 28-November-2013].

[Wikipedia2013b] Wikipedia. 2013b. Pearson product-moment correlation coefficient — Wikipedia, the free encyclopedia. [Online; accessed 28-November-2013].